Project 2 for Bayesian Inference
Spring 2009

You can use functions other than `runif(0,1)` in order to simulate from a suitable distribution.

1. Let $\pi(x, y)$ denote a bivariate density which is uniform over the region bounded by the lines $x = 0$, $x = 1$, $x + y = 1$, and $x + y = 2$.

   (a) Identify the conditional densities $\pi(y|x)$, $\pi(x|y)$ for $0 < x < 1$ and $0 < y < 2$. (A sketch of the region will help). Design and implement a Gibbs sample to sample from π, by drawing from these conditional densities.

   (b) Generate histogram estimates of the marginal densities $\pi(x)$ and $\pi(y)$, and verify that your algorithm produces estimates that agree with the theoretical marginal densities (which you will have to work out!)

   (c) Suppose that you used a similar algorithm to sample from a density that was uniform on the region bounded by $x = 0$, $x = 1$, $x + y = 1$, and $x + y = 1.1$. By considering the shape of this region, suggest why your Gibb's sampler would not explore the target density efficiently.

2. Design and implement a Gibbs sampler to simulate from the posterior density $\pi(\mu, \psi|\mathbf{x})$ where $\mathbf{x}$ is a random sample of size n from a Normal distribution with unknown mean mu, and variance $\sigma^2 = \frac{1}{\psi}$, using independent normal and Gamma priors for μ and ψ respectively. (Code for this has been given out to the class.)

   (a) Apply your algorithm to investigate $\pi(\mu, \psi|\mathbf{x})$ where $\mathbf{x}$ is a random sample of size 20 from an N(1, 0.5) density (which you will have to generate yourself! - Remember rnorm() uses the standard deviation.) Use initially the improper, vague priors ($\pi(\mu) \propto 1$, $\pi(\psi) \propto \psi^{-1}$) discussed in lectures. By generating a suitably long sequence of iterates from the chain, estimate the posterior mean and variance of the parameters, and the posterior probabilities that $\mu > 1.5$ and $\sigma^2 > 0.75$.

   (b) Repeat these calculations for a sample size of 60 from an N(1, 0.5) distribution.

(c) Consider how you might use standard results (see chapter 3 of notes!) in order to check whether your Gibbs sampler is working correctly.

3. Let $\mathbf{x} = (5.25, 4.80, 4.55, 5.8, 5.3, 4.38, 3.08, 5.60)$ denote a random sample from a $\Gamma(\alpha, \beta)$ distribution for which $\sum x_i = 38.76$ and $\prod x_i = 266, 274$. Assume that *a priori* $\alpha \sim U(1, 15)$ and $\beta \sim Exp(0.1)$. Implement in R the Metropolis algorithm described in lectures for simulating from the posterior $\pi(\alpha, \beta|\mathbf{x})$. By applying the method of moments to the data, identify suitable initial values for α and β. Use trial and error to identify suitable step-sizes for updates to α and β.

   (a) Estimate the posterior mean and variance of α and β from a suitably long run of the chain. Examine the shape of the marginal histogram for α and estimate its posterior mode.

   (b) Estimate 90% equal-tailed credible intervals for α and β from the output of the chain. (You can do this using the sort() command in R which arranges the elements in a vector in increasing order. The end points of a credible interval can then be obtained from the ordered Markov chain output.)

   (c) How do the marginal distributions of α and β change when the prior for β is selected to be a) $Exp(1)$ and b) $Exp(5)$?

   (d) By plotting the points (α_i, β_i) on a scatter diagram, investigate the dependence of α and β in the posterior distribution.

   (e) The data were generated from a $Gamma(8, 1.5)$ distribution. Simulate random samples of size 20 and size 40 from this distribution using the rgamma(n, alpha, beta) function in R and apply your algorithm to these samples in order to estimate (α, β) using the $Exp(0.1)$ prior for β. Investigate how the posterior marginal densities for the parameters change as the sample size becomes larger.

4. Modify your code for the M-H sampler of the previous question (inference on (α, β) in the $\Gamma(\alpha, \beta)$) distribution) by using a Gibbs update for β instead of the Metropolis step. By examining trace plots of the values of α and β against iterate and/or calculating autocorrelation functions determine whether the Gibbs sampler has superior mixing properties.